Time Synchronization Comparison of Network Interface Cards Under Load

Yury Kamen (presenter), Yakov Kamen, Ankur Sharma, Alex Wilms, Stan Chernavsky

Equinix Inc.

At a Glance

- We are presenting test results of time synchronization for common enterprise grade NICs located in a high performance colocation center.
- We collected and analyzed synchronization via NTP and PTP protocols on Linux, and NTP on Windows.
- We performed tests using a special Precise Timing Platform which integrates GPS-based Master Devices with cloud exchange infrastructure.

Problems we are trying to solve

- There are several expensive high-precision timing solutions that provide nanosecond level precision. What is needed by thousands of enterprise customers (worldwide) is a system that is secure, simple to operate, robust, precise, and inexpensive.
- Typical customers use generic enterprise quality NIC cards with HW timestamps for clock synchronization. Here we present test results that show how these cards behave under different network load.
- These devices are distributed worldwide and customers require different levels of accuracy, support of different operating systems, and time synchronization protocols. They use different NICs and operate under the certain network load.

Testing and Measurement Environment

- We use the following parameters:
  - Performance of time synchronization on LAN and WAN
  - OS: Linux CentOS 7 and Windows 10
  - Protocols: NTP (Linux, Windows) and PTP (Linux)
  - Network Load: 0%, 10%, 30%, 50%
  - Sync. Packet Rate (1/16s Linux PTP, 1s Linux NTP, 1s Windows NTP)
  - NIC: Mellanox ConnectX-3 Pro (CX-3 Pro) , Solarflare SFN5122F , Intel® Ethernet Controller I210-AT
  - HW: SuperServer 5019S-M, Dell OptiPlex 7050

Equinix Testing Platform Setup

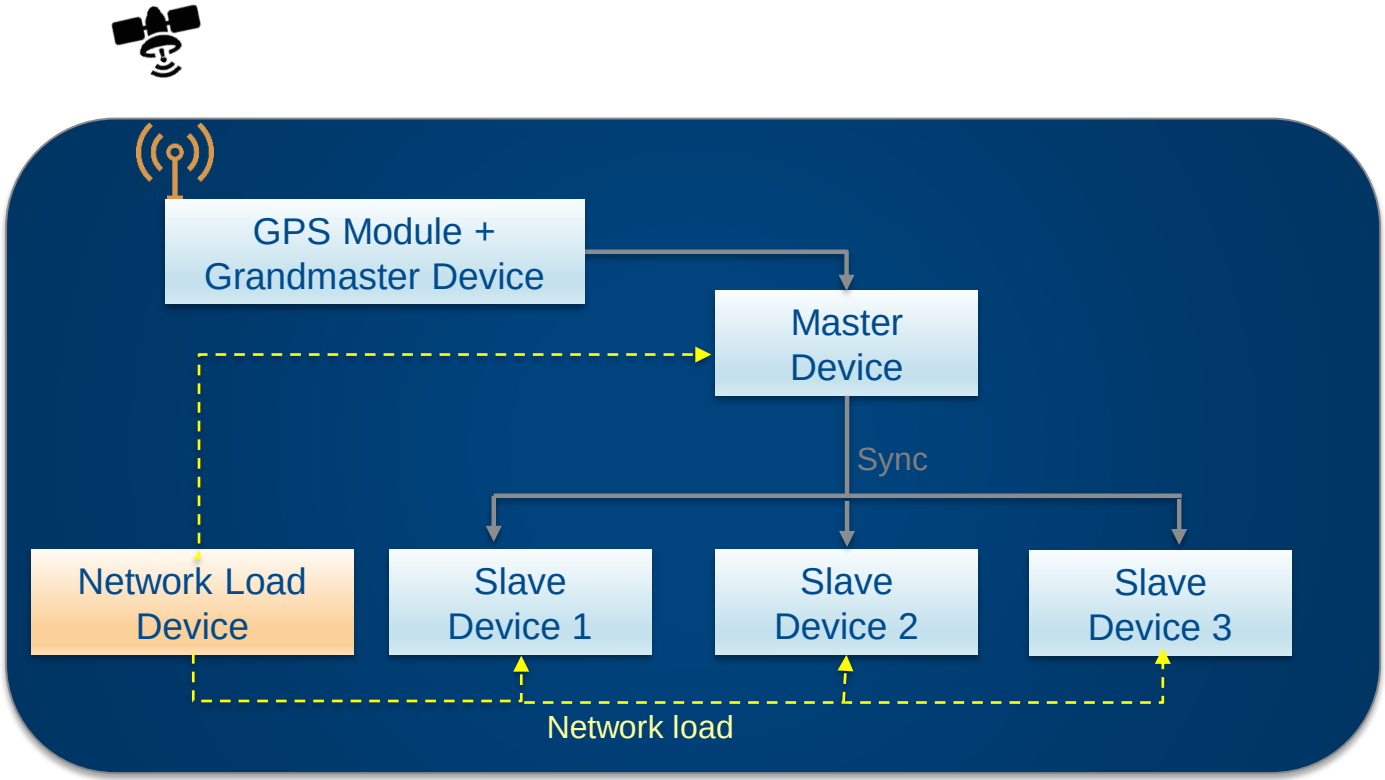

Questions we answer with testing:

- How does the load of different devices affect time synchronization precision in our setting?
- How does synchronization precision degrade when the network load increases?

EXPERIMENTS AND RESULTS

Examples of Tests Under Different Network Load

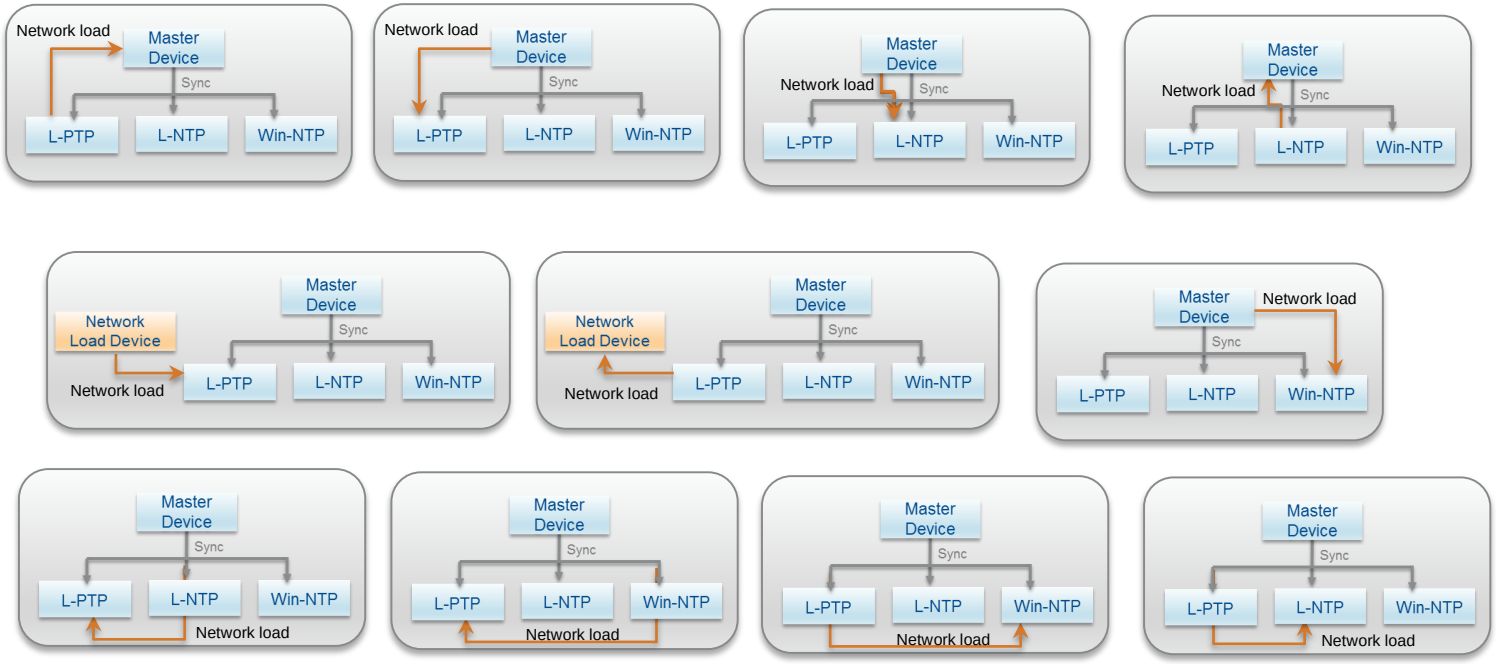

EQUINIX

Tested Network Cards

- Mellanox ConnectX-3 Pro (CX-3 Pro)
- Intel® Ethernet Controller I210-AT
- Solarflare SFN5122F

NIC Accuracy Comparison

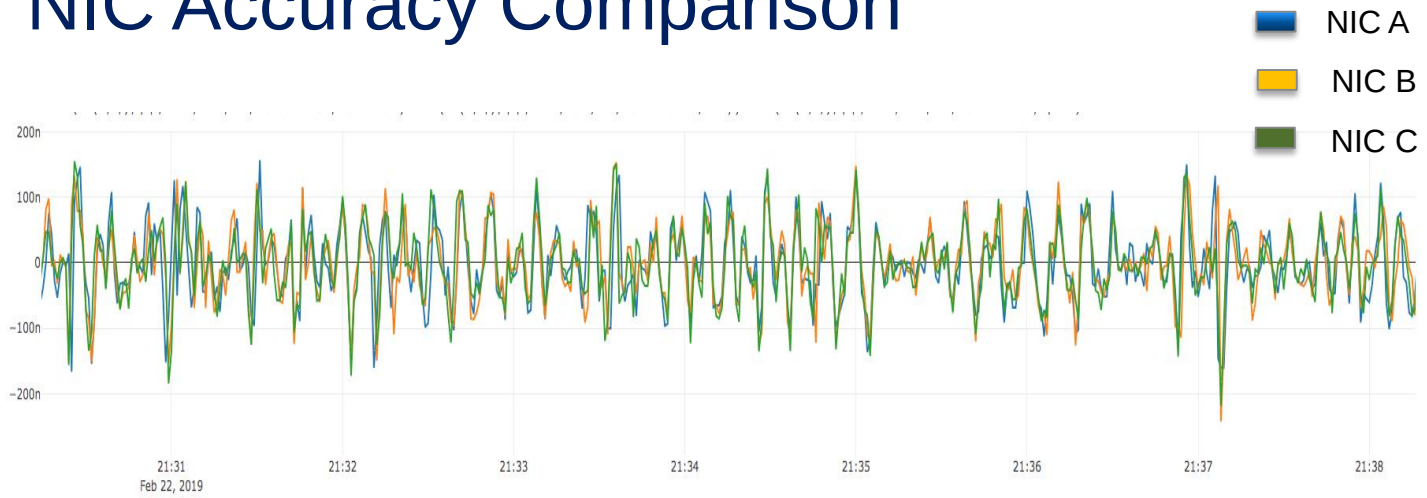

- Time offsets on all NICs follow the same pattern with both PTP and NTP.
- All NICs provide statistically indistinguishable accuracy for the 0-10% load (PTP, NTP).
- Under the load 10-30% STD of NICs fluctuate up to 30% (PTP).
- At load 50% and above some NIC accuracy significantly degrades (PTP).

Linux PTP: Statistics

Statistics	Network Load*			
	0%	10%	30%	50%
average offset (ns)	[-2:2]	[-11; 5]	[-44 ;27]	[-7; 36]
std (ns)	[14,54]	[26; 57]	[27; 185]	[29; 376]
max offset (ns)	137	265	650	829
min offset (ns)	-128	-141	-363	-685

Measured: Offset of the slave NIC clock from the master time source
Number of measurements: ~ 130,000
Intel 50% is excluded

Observations:

- Load plan does not significantly affect Linux PTP NIC synchronization
- Network load from 0% to 50% does not significantly affect average time offset
- Offset jitter steadily increases with network load
- With 50% network load Linux PTP NICs continue to provide sub-3µs accuracy. It's significantly less than 100µs MiFID II requirement for the EU financial markets.

EQUINIX

Time Sync Charts for Linux PTP

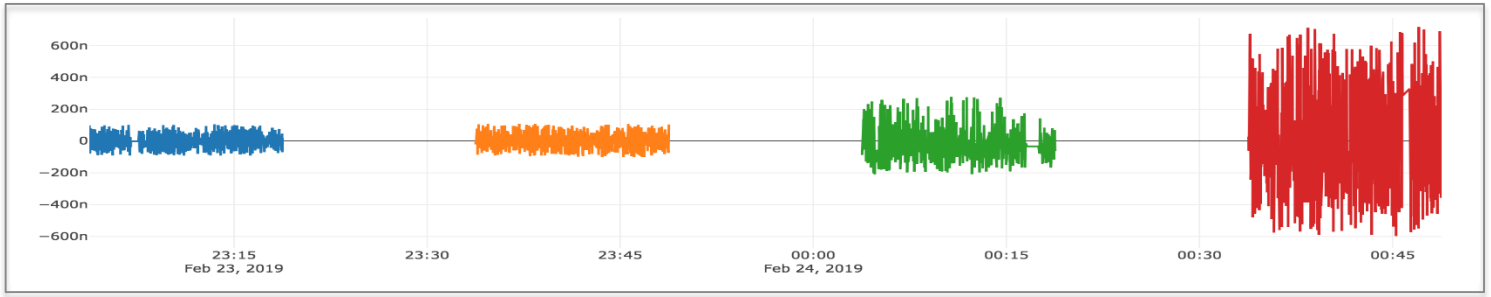

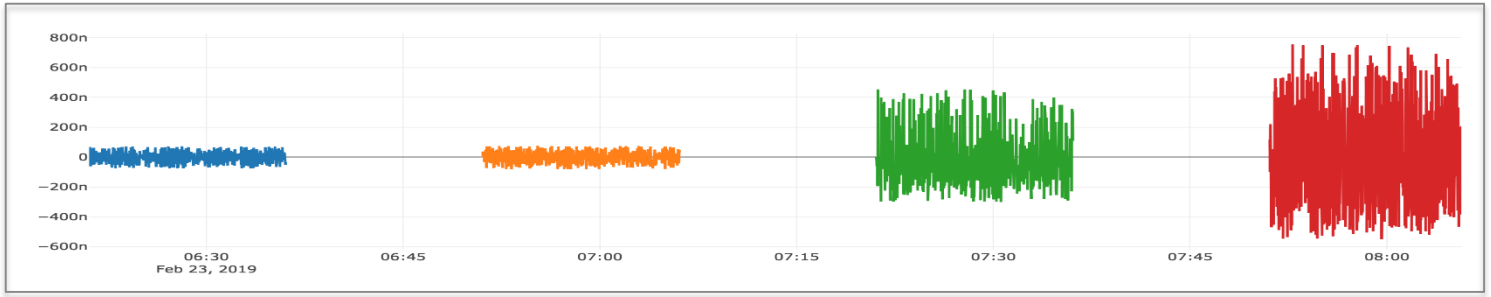

Load 1% Load 10% Load 30% Load 50%

EQUINIX

Linux NTP: Statistics

Statistics	Network Load			
	0%	10%	30%	50%
average offset (μs)	[-12.9;-0.7]	[-14.84; -3.36]	[-14.2;-0.8]	[-13.0; -1.1]
std (μs)	[0.4-9.5]	[2.03; 9.0]	[0.8; 9.8]	[0.9; 9.0]
max offset (μs)	0.1	0.0	1.0	2.5
min offset (μs)	-29.0	-31.0	-30	-31

number of measurements: ~ 75,000

Observations:

- Load plan does not significantly affect Linux NTP NIC synchronization;
- Offset jitter is similar for all loads
- With 50% network load Linux NTP NICs deliver sub 50us accuracy.

Time Sync Charts for Linux NTP

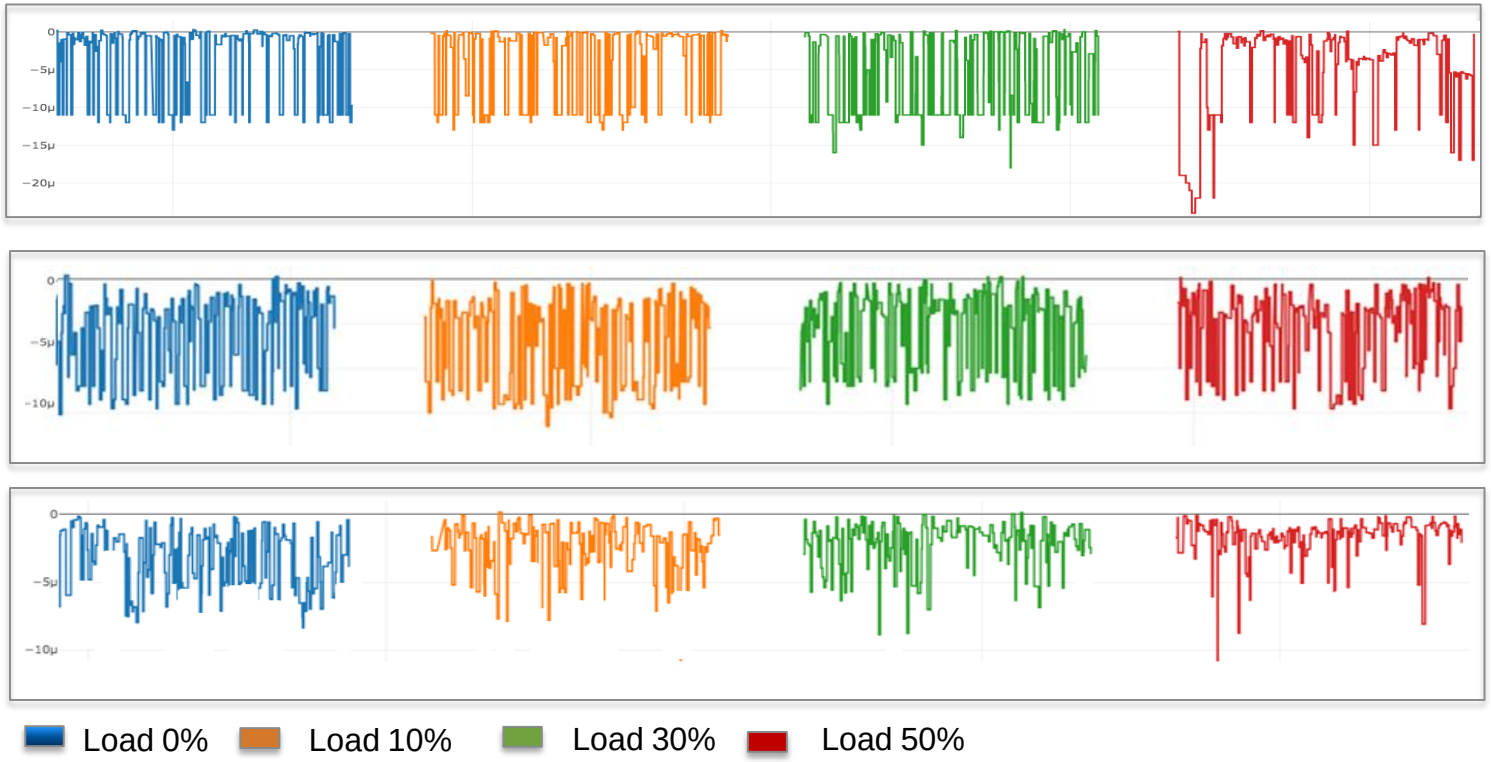

EQUINIX

Windows NTP

- Windows NTP does not use NIC HW timestamps
- Time synchronization accuracy does not depend on a NIC
- Network load below 50% does not affect Windows NTP performance
- LAN connected devices (Stratum 3) provide time synchronization accuracy under 2ms.

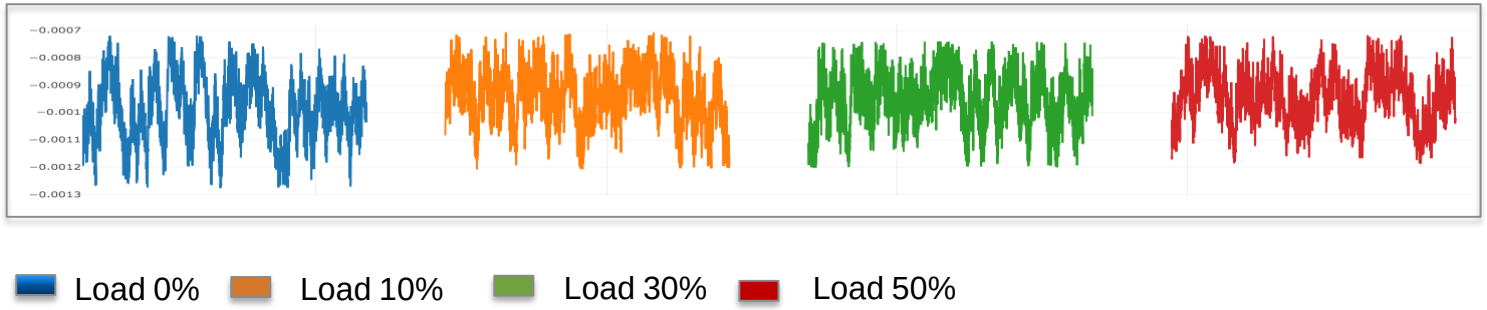

EQUINIX

Windows NTP Analytics

Stratum 3 (LAN) vs. Windows NTP Internet

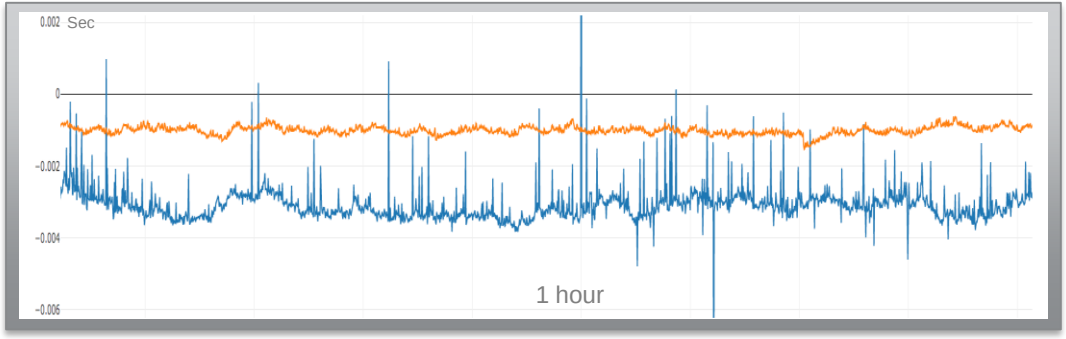

	Median	STD	Average	Max	Min
Win NTP Stratum 3	-0.96ms	0.22ms	-1.1ms	-0.2ms	-1.6 ms
Win NTP Internet	-2.9ms	2.06ms	33.5ms	4.6s	-9.5s

EQUINIX

VISUALIZATION AND DATA

ANALYTICS SYSTEM

EQUINIX

Visualization System

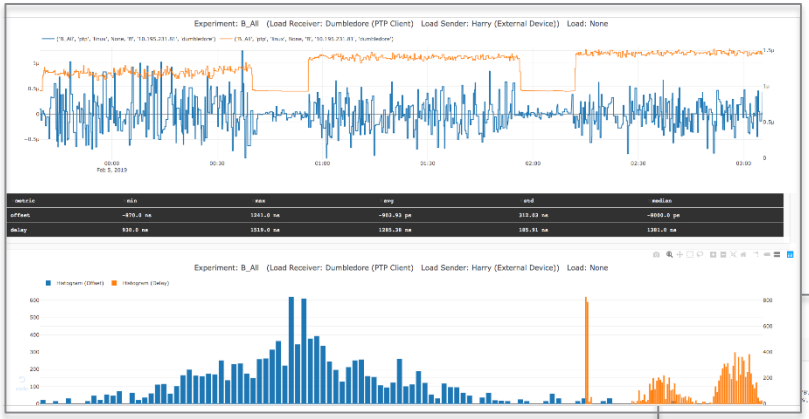

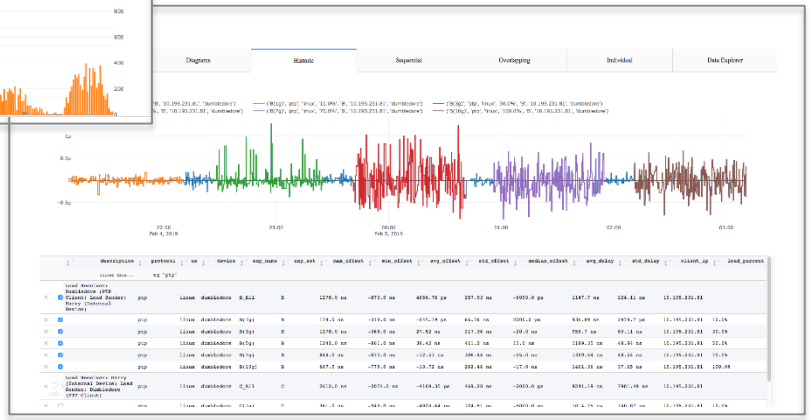

Features:

- Interactive Charts
- Dynamic calculation and display of statistics and histograms
- Network load visualization

EQUINIX

Conclusions: Linux NTP/PTP

- All tested NICs demonstrated respectable level of accuracy. When network load is below 50% Linux PTP-based synchronization delivers sub 3ms accuracy.
- NTP protocol provides sub 100ms precision with 50% network load. It shows ~30x lower accuracy than PTP protocol with 50% load.
- All tested cards demonstrated similar performance under load
- No hardware compatibility issues were found

Conclusions: Windows NTP

- Windows NTP implementation does not use HW timestamps.
- LAN Stratum 3 Windows NTP device provides sub 2ms accuracy.

Work in Progress

- Our testing process is in progress. We are planning to test NIC behavior under dynamically changing load.
- We are planning to test Windows PTP with both software and hardware timestamping modes.

Collaborative Testing

- We are open to collaborate with the community to test additional NICs.
- Please contact us if you are interested to test NICs and/or HW systems.

Contact Information

- ykamen@equinix.com
- kyakov@equinix.com

EQUINIX